



**NAVAL
POSTGRADUATE
SCHOOL**

MONTEREY, CALIFORNIA

**NAVAL INTEGRATION INTO
JOINT DATA STRATEGIES AND ARCHITECTURES IN JADC2**

by

Arkady Godin

October 2022

Approved for Public Release; Distribution is unlimited.

Prepared for: OPNAVN2N6F33.

This research is supported by funding from the Naval Postgraduate School,
Naval Research Program (PE 0605853N/2098).

NRP Project ID: NPS-22-N279-A

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.

1. REPORT DATE 10/22/2022		2. REPORT TYPE Technical Report		3. DATES COVERED START DATE 10/24/2021 END DATE 10/22/2022	
3. TITLE AND SUBTITLE Naval Integration into Joint Data Strategies and Architectures in JADC2					
5a. CONTRACT NUMBER		5b. GRANT NUMBER		5c. PROGRAM ELEMENT NUMBER 0605853N/2098	
5d. PROJECT NUMBER NPS-22-N279-A, W2223		5e. TASK NUMBER		5f. WORK UNIT NUMBER	
6. AUTHOR(S) Arkady Godin					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School, 1 University Circle, Monterey CA 93943-5000				8. PERFORMING ORGANIZATION REPORT NUMBER NPS-IS-22-003	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) OPNAV N2N6F33, 1200 Navy Pentagon, Washington DC				10. SPONSOR/MONITOR'S ACRONYM(S) NRP; OPNAV	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) NPS-IS-22-003; NRP NPS-22-N279-A	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUBJECT TERMS Joint, Command and Control, Knowledge Strategy, Situation, Context, Summarization, Joint Organizations, Shared Missions, Collaboration, Knowledge Representation, Logics, Neuro-Symbolic Approach, Machine Learning, Artificial Intelligence, Symbols, Learning, Reasoning					
14. ABSTRACT Joint All Domain Command and Control (JADC2) is not a new problem. Each US service and intelligence agency has been fighting as joint in the past and in the present. It's just there was never a challenge that integration between organization had to be using uniform means of integration as opposed to proprietary ones as defined by organization leading the joint mission. As usual, there is always a need for highly horizontal integration. What seems to be new is to ensure vertical integration between hierarchically organized roles in each of the tiers of JADC2 scales greatly via elegant and highly effective design. This needed not only due to a joint cross-organizational structure. Due to the battlespace complexity, there is a critical need for each role of the command to be easily substitutable in case of casualties including reduction in human readiness.					
16. SECURITY CLASSIFICATION OF:				17. LIMITATION OF ABSTRACT UU	
1. REPORT Unclassified	2. ABSTRACT Unclassified	C. THIS PAGE Unclassified		18. NUMBER OF PAGES 52	
19a. NAME OF RESPONSIBLE PERSON Arkady Godin				19b. PHONE NUMBER (Include area code) (831) 224-5699 cell	

THIS PAGE INTENTIONALLY LEFT BLANK

**NAVAL POSTGRADUATE SCHOOL
Monterey, California 93943-5000**

Ann E. Rondeau
President

Scott Gartner
Provost

The report entitled “Naval Integration into Joint Data Strategies and Architectures in JADC2” was prepared for OPNAV N2N6F33 and funded by the Naval Postgraduate School, Naval Research Program (PE 0605853N/2098).

Further distribution of all or part of this report is authorized.

This report was prepared by:

Arkady Alexander Godin
Faculty Associate - Research

Reviewed by:

Released by:

Alex Bordetsky, Chairman
Information Sciences Department

Kevin B. Smith
Vice Provost for Research

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

ACKNOWLEDGEMENT	VIII
INTRODUCTION <i>SHOULD ANY DATA STRATEGY BE JOINT?</i>	10
PROBLEM STATEMENT <i>IS IT KNOWLEDGE REPRESENTATION PROBLEM?</i>	12
EXECUTIVE SUMMARY	14
CHAPTER 1 <i>SURVEYING KNOWLEDGE REPRESENTATIONS</i>	17
Nonmonotonic Reasoning	17
The frame problem	18
Answer Sets	19
Reasoning in Dynamic Domains	20
Qualitative Modeling	21
Bayesian Networks	22
Temporal Representation and Reasoning	22
Qualitative Spatial Representation and Reasoning	23
Reasoning about Knowledge and Belief	24
Situation Calculus	25
Event Calculus	26
Temporal Action Logics	26
Nonmonotonic Causal Logic	27
CHAPTER 2 <i>NEURO-SYMBOLIC LEARNING AND REASONING</i>	28
CoLlision Events for Video REpresentation and Reasoning (CLEVRER)	29
Approach of various state-of-the-art visual reasoning models on CLEVRER	29
Conclusion by evaluating team of CLEVRER with 4 types of questions	29
Modular enhancements to neuro-symbolic dynamic reasoning model	30
CHAPTER 3 <i>“COGENT WAY” KNOWLEDGE STRATEGY</i>	32
CHAPTER 4 <i>NPS/NAVWAR KNOWLEDGE STRATEGY EXTENSION</i>	33
Battlespace Situations and Context for Battlespace Situation	33
Summarization of Situations and Situational Context	35
Embracing 3 rd Wave of AI: Contextual Adaptation	36
Moving from Metadata Catalog to Contextual Data Catalog	38
Stepping into the World of Contextual Data Catalog	39
CHAPTER 5 CONCLUSIONS	40
LIST OF ACRONYMS	41
LIST OF REFERENCES	43
INITIAL DISTRIBUTION LIST	49

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF FIGURES

Figure 1	From Representation to Evaluation to Optimization	13
Figure 2	Data is Coal and Knowledge is Gold	38
Figure 3	Metadata Catalog (single track) vs Contextual Data Catalog	39

ACKNOWLEDGEMENT

My sponsor, Bill Treadway, OPNAV N2/N6, was, as usual, quite enthusiastic while demanding, to keep the study on the right track. Mr. Treadway's goals have always been to express recommendations in a simple to grasp language to ensure both the sponsor and his upper leadership are clear on the investments the study is recommending. My alter ego Rich Chase, NAVWAR, Data Strategy Warrant Holder had been always providing me with as much time as I needed to ensure my ideas resonate with him. Rich was ensuring they are solid and address the real warfighting needs. I also want to thank Mike Green, my colleague at the Naval Postgraduate School. Mike was always inspirational that my abstract approach to the data strategy for Naval and Joint Staff All-Domain Command & Control is exactly what is necessary to construct the right integration-ready knowledge strategy core. Big thanks to my other partners on the faculty at the Naval Postgraduate School, especially Scot Miller, Navy CAPT (Retired), and Dr. Bonnie Johnson.

THIS PAGE IS INTENTIONALLY LEFT BLANK

INTRODUCTION

SHOULD ANY DATA STRATEGY BE JOINT?

The Data Strategy designed by NRP FY20 Broad Area Study [1] has been designed for the purpose on “describe how the Navy will analyze and transmit data in the future distributed fleet” [2]. The author of [1] also stipulated “*Since information and knowledge quality depend so heavily on data management, the latter is viewed by DoD innovators as the cornerstone for the future digital mission data architecture in the Joint Forces battlespace, since information and knowledge quality depend so heavily on it.*” In the final sub-chapter in [1], “Areas of the Future Research”, the author concluded that “... *executing this strategy will enable the Navy to join whatever Joint data is formulated in support of JADC2 and other current concepts.*”

This study is an opportunity to further advance an idea that earlier proposed Naval Data Strategy is, indeed, fits “joinable” requirements. However, with regards to “data management” strategy offered by JS J6 for JADC2, our team believes the time is right to look at “data” with strategic perspective in mind. We view JS J6 to work with the best thinkers to evolve “data management” strategy to support the requirements prescribed by the past and earlier Data Strategy efforts. We firmly believe “data management” strategy should take an important, but still a backseat by opening a stage to a knowledge strategy. We believe the latter not to be cast in stone but will be an alive mechanism providing warfighters at the tactical edge and decision making supporting their operations with knowledge-understanding capabilities to prevail in the joint battlespace by outmaneuvering and outsmarting adversaries in overly complex and uncertain decisioning space.

The final thought we want to share in the “Introduction” section of current Broad Area Study is to recognize that any Data Strategy must include provisions to be a “joint” one. Prior to the invention of JADC2 each service and agency had been and still continuing to operate as a “joint” force. Naval Data Strategy is “joint” organically as Navy and Marines are, with few possible exceptions, operate in a joint manner. However, Navy had been on joint missions shared with all other Services like Army and Airforce. The same applies to Navy Seals who are coordinated their efforts with other branches of the Special Operations Forces (SOF). It goes without saying that Naval service is tightly coupled with Space Forces, and such intelligence agencies as NSA, NRO and NGA.

THIS PAGE INTENTIONALLY LEFT BLANK

PROBLEM STATEMENT

IS IT KNOWLEDGE REPRESENTATION PROBLEM?

“Naval Integration into Joint Data Strategies in JADC2” broad area study topic assumes the data strategy for Naval and Joint Staff’s domains are able to support the concept of “jointness”. Otherwise, if naval data strategy is not capable to support the concept of “jointness”, it would be impossible for it to be of a benefit to a joint community at any tier: tactical, operational, or strategic.

The concept of “jointness”, essentially, adds more organizational dimensions to “joint cubes” abstraction where each cube is built over dimensions which represent organizational entities involved in a particular mission or set of missions. It is assumed that mission or multi-missions are necessary for conducting an operation.

What is a difference between naval and joint all-domain command & control data strategies? It is, simply, in a greater number of commands like services, intel, agencies, coalition partners, and any other entities required for inclusion into JADC2 and, naturally, a greater number of domains supported within each of the commands.

What is then the real problem for a “jointness” challenge? Each reader, likely, has a good understanding what entails to support Machine Learning and Deep Learning (ML/DL) paradigm. The answer revolves around whatever is necessary to develop the models from the data training and algorithmic point of view. Considering our focus being on a data side, a focus then is on organization for the data to get trained. What is the impact of such focus on “jointness” in JADC2? The answer is rather unspectacularly brief. ML/DL are looking for a data to generate information. Any data must be validated and labeled. Different data streams might require fusion. Such requirement is not unique to “Joint” type of analysis. It is there for non-joint SoSs as well.

A presented hypothesis what a primary challenge is to decide what kind of data is critical for the success of JADC2 joint operations. Our philosophy is that the data which is paramount for such success is not “raw data” from “sensored environments” but rather continuously derived knowledge and understanding. It means **“joint crown jewel”** data strategy requirement is to have “joint” common shared representation/reasoning structures/methods to store/derive knowledge and understanding.

Pedro Domingos in [99] organized all Five Tribes of the Artificial Intelligence the following way:




Figure 1 From Representation to Evaluation to Optimization

Professor Pedro Domingos has suggested using canonical knowledge representation based on Markov Logic Network (MLN) to represent any knowledge for all five tribes of AI including Symbolists, Connectionists, Evolutionaries, Bayesians and Analogizers. With the advent of Causal tribe, MLN can represent “directed causal graphs” as well. Sharing knowledge via MLN-enabled container opens a new chapter on merging knowledge across all domains in Command & Control. JADC2 knowledge integration is straightforward if compared with the data and storage integration.

EXECUTIVE SUMMARY

Joint All Domain Command and Control (JADC2) is not a new problem. Each US service and intelligence agency has been fighting as joint in the past and in the present. It's just there was never a challenge that integration among organizations had to be using uniform means of integration as opposed to proprietary ones as defined by organization leading the joint mission. As usual, there is always a need for highly horizontal integration. What seems to be new is to ensure vertical integration between hierarchically organized roles in each of the tiers of JADC2 scales greatly via elegant and highly effective design. This is needed not only due to a joint cross-organizational structure. Due to the battlespace complexity, there is a critical need for each role of the command to be easily substitutable in case of casualties including reduction in human readiness.

We are convinced that Situation Awareness (SA) rooted in the knowledge of the evolution of each situation, from the point of origination to the point of deprecation, is essential. It is a must prerequisite for supporting significant volatility of the situations in the modern battlespace. Knowing situations ensures our knowledge of the battlespace is both contextual and adaptive to the dynamic ad-hoc natures of the world. We must continuously account for not only relations between situations' participants (subjects and objects) and actions which change the states of the world model's state machine. It is imperative to realize the environments we are fighting in, such as METOC, EMW, Acoustic, are highly dynamic, resulting in changing conditions, and effects in the form of environmental events. These external threshold events contribute to pro-active generation of external factors resulting in new conditions due to continuously changing causes to the situations and ongoing planning and replanning activities. We are confident concurrent environments play significant roles in the formation of mission contexts under continuously evolving situations.

We believe designing the enterprise bottom-up from the tactical edge is essential due to the knowledge of situations and contexts which may be obtained exclusively from the tactical edge. Sharing contextual knowledge with higher roles at the tactical tier and tiers above (i.e., operational, and strategic/national, is the only strategy to manage SA to ensure a comprehension of currently assessed and projected situations. Our team is firmly believing that imaginary 'what-if' questions are critical as they reveal causal relations by uncovering counterfactuals. In fact, "counterfactuals" is a mechanism to detect novel situations which is a necessary prerequisite to situations' detection.

We are confident our study is making significant contribution to decision-making actors at different roles across tactical, operational, and strategic/national tiers. Our conviction is that decision-ready knowledge must be delivered to all decision-making actors without any further transformation. This knowledge will be in context at a proper level of granularities for every decision-maker. This automatic push-driven streaming delivery will ensure the timeliness of providing valuable knowledge with an opportunity to decision to ask for me by querying the back end based on the provided contexts equipped with the knowledge of situations.

Making knowledge delivery requires back end understanding on what decision makers view as valuable knowledge for their respective Areas of Responsibilities (AORs). There is also a need to support a concept of “knowledge in-situ” for extracted knowledge at the tactical edge. This requires imbedding collected observations and persisted novel detected situations into intelligent information-knowledge structures suitable for reasoning by various logics including causal, situation, action, event and any others, if necessary. If “smart data-information structures” is employed moving information-knowledge will be limited to a bare minimum which is of paramount importance for limited network bandwidth at the edge.

Considering a need to summarize captured information and knowledge is beneficial to a multitude of decision-makers, there are couple of preconditions to realize it. First of all, intelligent information-data structures must be embedded into the “summarization containers” designed to summarize graphs, other structures capable of capturing smart structures, events and world model. Secondly, all “smart information-knowledge structures” must support composability. Making it happen requires employing ontological “type system” to types are representable as compositions.

Presented executive summary touches upon a variety of interdisciplinary concepts. This is a necessity as information and knowledge for the logical world model is highly heterogenous, multi-dimensional and hierarchical. “Type system” is a foundational concept that is absolutely necessary to represent complex ever-changing world model. We are convinced JADC2 requires employment of the summarization aggregation engine to feed knowledge to all decision-makers.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 1

SURVEYING KNOWLEDGE REPRESENTATIONS

Knowledge Representation and Reasoning (KRR) is pursued by a group of AI scientists doing research on classical AI focused on symbolic reasoning. These scientists offer a prove that symbolic reasoning is critical for decision-making enabled by AI. The authoritative, most-used AI textbooks, adopted by over 1500 schools, were written by John Sowa [2] and Peter Norvig [3].

Based on initial literature research, Knowledge Representation (KR) may be classified into the following categories:

- Logics-based Knowledge Representation
- Semantic Network
- Frames
- Rule-based Engine

Due to study's research focused on JADC2, research team finds Logics-based KRR as the most appropriate to JADC2 due to the ad-hoc nature of causally induced events, actions with deterministic and non-deterministic effects on missions' situations. Therefore, sections will be dedicated to Logics-based KRR.

Semantic Network has been discounted as it uses first-order logic predicate calculus provided non-advanced reasoning over entity relations in the graphs. Frames logic (F-Logic) stands in the same relationship to object-oriented programming (OOP) as classical relational calculus stands to relational database programming [4]. Research team views OOP is pursuing the goals that are not congruent with the goals and operational specifics of JADC2 described at the onset of this paragraph.

Nonmonotonic Reasoning

Classical logic is *monotonic* in the following sense: whenever a sentence A is a logical consequence of a set of sentences T , then A is also a consequence of an arbitrary superset of T . In other words, adding information never invalidates any conclusions.

Commonsense reasoning is different. We often draw plausible conclusions based on the assumption that the world in which we function and about which we reason is *normal* and *as expected*. This is far from being irrational. To the contrary, it is the best we can do in situations in which we have only incomplete information. However, as unexpected as it may be, it can happen

that our normality assumptions turn out to be wrong. New information can show that the situation actually is abnormal in some respect. In this case we may have to revise our conclusions.

For example, let us assume that Professor Jones likes to have a good espresso after lunch in a campus cafe. You need to talk to her about a grant proposal. It is about 1:00 pm and, under normal circumstances, Professor Jones sticks to her daily routine. Thus, you draw a plausible conclusion that she is presently enjoying her favorite drink. You decide to go to the cafe and meet her there. As you get near the student center, where the cafe is located, you see people streaming out of the building. One of them tells you about the fire alarm that just went off. The new piece of information invalidates the normality assumption and so the conclusion about the present location of Professor Jones, too.

Such reasoning, where additional information may invalidate conclusions, is called *nonmonotonic*. It has been a focus of extensive studies by the knowledge representation community since the early eighties of the last century. This interest was fueled by several fundamental challenges facing knowledge representation such as modeling and reasoning about rules with exceptions or *defaults* and solving the *frame* problem. [5]

The frame problem

To express effects of actions and reason about changes in the world they incur, one has to indicate under what circumstances a proposition whose truth value may vary, a *fluent*, holds. One of the most elegant formalisms to represent change in logic, *situation calculus*, uses situations corresponding to sequences of actions to achieve this. For instance, the fact that Fred is in the kitchen after walking there, starting in initial situation S_0 , is represented as

holds (in (Fred, Kitchen), do (walk (Fred, Kitchen), S_0)).

The predicate *holds* allows us to state that a fluent, here *in (Fred, Kitchen)*, holds in a particular situation. The expression *walk (Fred, Kitchen)* is an action, and the expression *do (walk (Fred, Kitchen), S_0)* is the situation after Fred walked to the kitchen, while in situation S_0 .

In situation calculus, effects of actions can easily be described. It is more problematic, however, to describe what does *not* change when an event occurs. For instance, the color of the kitchen, the position of chairs, and many other things remain unaffected by Fred walking to the kitchen. The frame problem asks how to represent the large number of non-changes when reasoning about action. [6]

Answer Sets

Answer Set Prolog is a language for knowledge representation and reasoning based on the *answer set/stable model* semantics of logic programs [7, 8]. The language has roots in declarative programming [9, 10], the syntax and semantics of standard Prolog [11], disjunctive [12, 13] and nonmonotonic logic [14, 15, 16]. Unlike “standard” Prolog it allows us to express disjunction and “classical” or “strong” negation. It differs from many other knowledge representation languages by its ability to represent *defaults*, i.e., statements of the form “*Elements of a class C normally satisfy property P*”. Answer Set Prolog provides a powerful logical model of this process. Its syntax allows for the simple representation of defaults and their exceptions, its consequence relation characterizes the corresponding set of valid conclusions, and its inference mechanisms often allow a program to find these conclusions in a reasonable amount of time.

There are other important types of statements which can be nicely expressed in Answer Set Prolog. This includes the causal effects of actions (“statement F becomes true as a result of performing an action a ”), statements expressing a lack of information (“it is not known if statement P is true or false”), various completeness assumptions “statements not entailed by the knowledge base are false”, etc.

The method of solving various combinatorial problems by reducing them to finding the answer sets of Answer Set Prolog programs which declaratively describe the problems is often called the *answer set programming paradigm (ASP)* [17, 18]. It has been used for finding solutions to a variety of programming tasks, ranging from building decision support systems for the Space Shuttle [19] and program configuration [20], to solving problems arising in bioinformatics [21], zoology and linguistics [22]. On the negative side, Answer Set Prolog in its current form is not adequate for reasoning with complex logical formulas—the things that classical logic is good at—and for reasoning with real numbers.

There is a substantial number of natural and mathematically elegant extensions of the original Answer Set Prolog. A long-standing problem of expanding answer set programming by aggregates—functions on sets—is approaching its final solution in [23, 24, 25, 26, 27]. The rules of the language are generalized [28] to allow nested logical connectives and various means to express preferences between answer sets [29, 30, 31]. Weak constraints and consistency restoring rules are introduced to deal with possible inconsistencies [32, 33]. The logical reasoning of Answer

Set Prolog is combined with probabilistic reasoning in [34] and with qualitative optimization in [35]. All of these languages have at least experimental implementations and an emerging theory and methodology of use.

Reasoning in Dynamic Domains

We assume that such a domain is modeled by a *transition diagram* with nodes corresponding to possible states of the domain, and arcs labeled by actions. An arc (σ_1, a, σ_2) indicates that execution of an action a in state σ_1 may result in the domain moving to the state σ_2 . If for every state σ_1 and action a , the diagram contains at most one arc (σ_1, a, σ_2) then the domain is called *deterministic*. The transition diagram contains all possible trajectories of the domain. Its particular history is given by a record of observations and actions. Due to the size of the diagram, the problem of finding its concise specification is not trivial and has been a subject of research for a comparatively long time. Its solution requires a good understanding of the nature of causal effects of actions in the presence of complex interrelations between fluents—propositions whose truth value may depend on the state of the domain. An additional level of complexity is added by the need to specify what is not changed by actions. The latter, known as the *frame problem* [36], is often reduced to the problem of finding a concise and accurate representation of the *inertia axiom*—a default which says that *things normally stay as they are*. The search for such a representation substantially influenced AI research during the last twenty years. An interesting account of history of this research together with some possible solutions can be found in [37].

There is also a substantial cross-fertilization between answer set based reasoning about actions and change and other similar formalisms including Situation Calculus [38, 39], Event Calculus [40, 41], and various temporal logics. There are, for instance, logic programming-based counterparts of Situation Calculus, which allow elegant solutions to the frame and ramification problem. Original versions of Event Calculus were directly expressed in the language of logic programming. The ability of temporal logic to reason about properties of paths is modeled by logic programming-based specification of goals in [42]. There is an example of the use of Answer Set Prolog and its reasoning methods for representing and reasoning about commonsense and linguistic knowledge needed for intelligent question answering from natural language texts. There are several interesting efforts of combining Answer Sets with Bayesian net based probabilistic reasoning, which substantially increases expressive power of both knowledge representation

languages and promises to lead to efficient algorithms for answering some forms of probabilistic queries. Finally, new results establishing some relationship between Description Logic and Answer Sets [43] may open the way for interesting applications of Answer Sets to Semantic Web. [44]

Qualitative Modeling

Qualitative modeling concerns representation and reasoning about continuous aspects of entities and systems in a symbolic, human-like manner. People who have never heard of differential equations successfully reason about the common-sense world of quantities, motion, space, and time. They do so often in circumstances offering little information, using the ability to characterize broad categories of outcomes to ascertain what might happen. For many tasks this is enough: Knowing that a valuable fragile object might be pushed off a table is sufficient reason to rearrange things so that it cannot happen. For other tasks, knowing the possible outcomes suggests further analyses, perhaps involving more detailed models. For example, an engineer designing a tea warmer must keep the tea at a drinkable temperature, while not allowing it to boil. Reasoning directly with qualitative models can capture important behavior patterns, automatically producing descriptions that are closer to the level of what people call insights about system behavior, making them useful for science, engineering, education and decision-support. Capturing the representational and reasoning capabilities that enable robust reasoning about continuous systems is the goal of qualitative modeling.

Qualitative modeling is today most commonly referred to in the literature as qualitative reasoning, but we use qualitative modeling here to emphasize that the representational work in this area shared equal importance with work on reasoning techniques per se. (As will be seen below, the tradeoffs in them are deeply intertwined.) Qualitative physics has often been used for research in this area as well, since understanding physical systems has been a central focus of much of the work in the area. However, this term has become less popular as the applicability of these ideas to areas such as finance, ecology, and natural language semantics have been explored.

Bayesian Networks

A Bayesian network is a tool for modeling and reasoning with uncertain beliefs. A Bayesian network consists of two parts: a qualitative component in the form of a directed acyclic graph (DAG), and a quantitative component in the form of conditional probabilities. Intuitively, the DAG of a Bayesian network explicates variables of interest (DAG nodes) and the direct influences among them (DAG edges). The conditional probabilities of a Bayesian network quantify the dependencies between variables and their parents in the DAG. Formally though, a Bayesian network is interpreted as specifying a unique probability distribution over its variables. Hence, the network can be viewed as a factored (compact) representation of an exponentially sized probability distribution. The power of Bayesian networks as a representational tool stems both from this ability to represent large probability distributions compactly, and the availability of inference algorithms that can answer queries about these distributions without necessarily constructing them explicitly.

There are two interpretations of a Bayesian network structure, a standard interpretation in terms of probabilistic independence and a stronger interpretation in terms of causality. According to the stronger interpretation, the Bayesian network specifies a family of probability distributions, each resulting from applying an intervention to the situation of interest. These causal Bayesian networks lead to additional types of queries and require more specialized algorithms for computing them. The directed nature of Bayesian networks can be used to provide causal semantics for these networks, based on the notion of *intervention* [45], leading to models that not only represent probability distributions, but also permit one to induce new probability distributions that result from intervention. In particular, a *causal network*, intuitively speaking, is a Bayesian network with the added property that the parents of each node are its direct causes. For example, $Cold \rightarrow HeadAche$ is a causal network whereas $HeadAche \rightarrow Cold$ is not, even though both networks are equally capable of representing any joint distribution on the two variables. [46]

Temporal Representation and Reasoning

This section is about representing *knowledge* in all its various forms. Yet, whatever phenomenon we aim to represent, be it natural, computational, or abstract, it is unlikely to be static. The natural world is always decaying or evolving. Thus, computational processes, by their nature, are dynamic, and most abstract notions, if they are to be useful, are likely to incorporate change. Consequently, the notion of representations *changing through time* is vital. And so, we need a clear way of

representing both our temporal basis, and the way in which entities change over time. This is exactly what this chapter is about.

We aim to provide an overview of many of the ways temporal phenomena can be *modelled*, *described*, *reasoned about*, and *applied*. In this, we will often overlap with other chapters in this collection. Some of these topics we will refer to very little, as they will be covered directly by other chapters, for example, *temporal action logic* [47], *situation calculus* [48], *event calculus* [49], *qualitative spatial representation and reasoning* [50], *temporal constraint programming* [51], *automated planning* [47, 52], and *qualitative modeling* [53]. Other topics will be described in this chapter at a high-level:

- *automated reasoning*,
- *description logics*, in [54]; and
- *natural language*

The topics, such as *reasoning about knowledge and belief* [55], *query answering* [56] and *multi-agent systems* [57], will only be referred to very briefly.

This section is not intended to be a comprehensive survey of *all* approaches to temporal representation and reasoning, it does outline many of the most prominent ones, though necessarily at a high-level. If more detail is required, many references are provided. Indeed, the first volume of the *Foundations of Artificial Intelligence* series, in which this collection appears, contains much more detail on the use of temporal reasoning in Artificial Intelligence [58] while [59, 60, 61, 62, 63] all provide an alternative logic-based view of temporal logics. [64]

Qualitative Spatial Representation and Reasoning

Space is multidimensional and is not in general adequately represented by a single scalar quantity. Early attempts at qualitative spatial reasoning within the QR community led to the ‘poverty conjecture’ [65]. Although purely qualitative representations were quite successful in reasoning about many physical systems [66], there was much less success in developing purely qualitative reasoners about spatial and kinematic mechanisms and the poverty conjecture is that this is in fact impossible—there is no purely qualitative spatial reasoning mechanism. [50]

‘Poverty conjecture’ suggests that spatial representations consist of two parts: a *metric diagram*, which includes quantitative information and thus provides a substrate that can support perceptual-like processing, and a *place vocabulary*, which makes explicit qualitative distinctions

in shape and space relevant to the current task. The metric diagram can use floating-point numbers, algebra, or even high-precision arrays – whatever it uses, there must be enough detail to support answering spatial queries by calculation, and it must be capable of supporting the construction of place vocabularies. Place vocabularies consist of *places*, contiguous regions of space where some important property is constant. Computing the place vocabulary according to the needs of the problem ensures that the relevant distinctions are made. Defining the places in terms of elements in the metric diagram makes the diagram a good communication medium for diverse representations.¹

We pay attention in this section specifically to qualitative spatial, and spatio-temporal reasoning (henceforth QSR). The challenge of QSR then is to provide calculi which allow a machine to represent and reason with spatial entities without resort to the traditional quantitative techniques prevalent in, for, e.g., the computer graphics or computer vision communities.

There has been an increasing amount of research in recent years which tends to refute, or at least weaken the ‘poverty conjecture’. Qualitative spatial representations addressing many different aspects of space including topology, orientation, shape, size and distance have been put forward. [50]

Reasoning about Knowledge and Belief

An agent operating in a complex environment can benefit from adapting its behavior to the situation at hand. The agent’s choice of actions at any point in time can, however, be based only on its local knowledge and beliefs. When many agents are present, the success of one’s agent’s actions will typically depend on the actions of the other agents. These, in turn, are based on the other agents’ own knowledge and beliefs. It follows that to operate effectively in a setting containing other agents, an agent must, in addition to its knowledge about the physical features of the outside world, consider its knowledge about another agent’s knowledge. This line of reasoning can be extended to justify the need for using deeper levels of knowledge, of course. Moreover, the task of obtaining relevant knowledge and that of affecting the knowledge of other agents, become important goals in many applications. This crucial connection between knowledge and action is what makes knowledge and belief two of the most frequently used notions in everyday discourse.

¹ <https://www.qrg.northwestern.edu/ideas/qsidea.htm>

It also suggests that rigorous frameworks for reasoning about knowledge and belief can be of value when analyzing scenarios involving multiple agents. [55]

The topic of reasoning about knowledge and belief is described in detail in the book *Reasoning about Knowledge* [67].

Situation Calculus

The situation calculus is a logical language for representing changes. It was first introduced by McCarthy in 1963, [68] and described in further details by McCarthy and Hayes [69] in 1969.

The basic concepts in the situation calculus are *situations*, *actions* and *fluents*. Briefly, actions are what make the dynamic world change from one situation to another when performed by agents. Fluents are situation-dependent functions used to describe the effects of actions. There are two kinds of them, *relational* fluents and *functional* fluents. The former has only two values: true or false, while the latter can take a range of values. For instance, one may have a relational fluent called “*handempty*” which is true in a situation if the robot’s hand is not holding anything. We may need a relation like this in a robot domain. One may also have a functional fluent called *battery-level* whose value in a situation is an integer between 0 and 100 denoting the total battery power remaining on one’s laptop computer.

According to McCarthy and Hayes [69], a situation is “the complete state of the universe at an instance of time”. But for Reiter [70], a situation is the same as its history which is the finite sequence of actions that has been performed since the initial situation S_0 . We shall discuss Reiter’s foundational axioms that make this precise later. Whatever the interpretation, the unique feature of the situation calculus is that situations are first-order objects that can be quantified over. This is what makes the situation calculus a powerful formalism for representing change and distinguishes it from other formalisms such as dynamic logic [71].

To describe a dynamic domain in the situation calculus, one has to decide on the set of actions available for the agents to perform, and the set of fluents needed to describe the changes these actions will have on the world. For example, consider the classic blocks world where some blocks of equal size can be arranged into a set of towers on a table. The set of actions in this domain depends on what the imaginary agent can do. [48]

Event Calculus

The event calculus [72, 73, 74, 37, 75] is a formalism for reasoning about action and change. Like the situation calculus, the event calculus has actions, which are called *events*, and time-varying properties or *fluents*. In the situation calculus, performing an action in a situation gives rise to a successor situation. Situation calculus actions are hypothetical, and time is tree-like. In the event calculus, there is a single timeline on which actual events occur.

A *narrative* is a possibly incomplete specification of a set of actual event occurrences [76, 37]. The event calculus is narrative-based, unlike the standard situation calculus in which an exact sequence of hypothetical actions is represented.

Like the situation calculus, the event calculus supports context-sensitive effects of events, indirect effects, action preconditions, and the commonsense law of inertia. Certain phenomena are addressed more naturally in the event calculus, including concurrent events, continuous time, continuous change, events with duration, nondeterministic effects, partially ordered events, and triggered events. [49]

Temporal Action Logics

The study of frameworks and formalisms for reasoning about action and change [77, 78, 79, 37, 80, 81, 41] has been central to the knowledge representation field almost from the inception of Artificial Intelligence as a general field of research [82, 83].

The phrase “Temporal Action Logics” represents a class of logics for reasoning about action and change that evolved from Sandewall’s book on *Features and Fluents* [79] and owes much to this ambitious project. There are essentially three major parts to Sandewall’s work. He first developed a narrative-based logical framework for specifying agent behavior in terms of action scenarios. The logical framework is state-based and uses explicit time structures. He then developed a formal framework for assessing the correctness (soundness and completeness) of logics for reasoning about action and change relative to a set of well-defined intended conclusions, where reasoning problems were classified according to their ontological or epistemological characteristics. Finally, he proposed a number of logics defined semantically in terms of definitions of preferential entailment² and assessed their correctness using his assessment framework.

² Preferential entailment reduces the set of classical models of a theory by only retaining those models that are minimal according to a given preference relation, a strict partial order over logical interpretations [84]

Several of these logics were intended to correspond directly to existing logics of action and change proposed by others at the time, while the rest were new and were intended to characterize broad classes of reasoning problems which subsumed some of the existing approaches. Each of these definitions of preferential entailment were then analyzed using the assessment framework, giving upper and lower bounds in terms of the classes of reasoning problems for which they produced exactly the intended conclusions. Much insight was gained both in terms of advantages and limitations of previously proposed logics of action and change and in how one might go about proposing new logics of action and change in a principled manner with formal assessments included. [47]

Nonmonotonic Causal Logic

In the last 15 years, there have been *many* reasoning about action proposals incorporating more explicitly causal notions. The nonmonotonic causal logic described in this chapter was introduced in [85]. The most relevant prior work appears in [86, 87, 88]. A much fuller account of causal theories was published in 2004 [89].

An implementation of causal theories—the Causal Calculator (CCALC)—is publicly-available, and many of the above-cited papers describe applications of it. The key to this implementation is an easy reduction from (a subclass of) causal theories to classical propositional logic, by a method introduced in [85]. Thus, automated reasoning about causal theories can be carried out via standard satisfiability solvers. (The initial version of CCALC was due primarily to Norm McCain and is described in his PhD thesis [90]. Since then, it has been maintained and developed by Vladimir Lifschitz and his students at the University of Texas at Austin.)

CHAPTER 2

NEURO-SYMBOLIC LEARNING AND REASONING

Neuro-Symbolic field is focused on seamless unification of AI and ML to eliminate all the Cons of AI and ML to stay with Pros only for unified AI and ML. This field is actively researched since 2020. There are many ideas. Dr. Arabshahi, as neuro-symbolic algorithms developer, describes several research interests [91] which neural-symbolic army of researchers is tackling with. We find **Interpretability** as a critical area of research.

Interpretability: Despite the data-driven algorithms' success, current machine learning models struggle to extract commonsense knowledge from data alone. This is because data contains little information about the commonsense knowledge that went into labeling or annotating it. On the other hand, model-driven algorithms (e.g., rule-based systems) that are programmed for a specific task, explicitly represent commonsense knowledge in terms of interpretable rules. But these models often lack coverage and are susceptible to uncertainty. To use the best of both worlds, I develop novel Neuro-Symbolic learning algorithms, which are hybrid models that leverage the robustness of connectionist methods and the soundness of symbolic reasoning to effectively integrate learning and reasoning.

Dr. Forough Arabshahi envisions that the major benefit of neuro-symbolic learning and reasoning is in the rapid ability to extract and generate commonsense knowledge via the hybrid neuro-symbolic models. Discovered insights are made available, through a feedback loop, back to neuro-symbolic model for pretrained language models (LMs.)

ATOMIC-2020 is a new general-purpose commonsense knowledge graph (CSKG) that is not readily available in pretrained LMs. ATOMIC20-20 is a novel commonsense knowledge graph containing tuples whose relations are specifically selected to be challenging for pretrained language models to express. Empirical studies demonstrate that ATOMIC20-20 contains high-accuracy knowledge tuples across multiple novel relations not found in existing CSKGs or expressible by Language Models. ATOMIC20-20 can be effectively used as a training set for adapting language models as knowledge models to generate high quality tuples on-demand. [93]. It is not surprising why Dr. Yejin Choi asked Dr. Forough Arabshahi to present her research findings to Allen Institute for AI [92]. It is likely that integrating neuro-symbolic algorithms developed by a team led by Dr. Arabshahi into ATOMIC-2020 will lead to discovering new common sense novel knowledge for ATOMIC-2020 CSKG.

CoLlision Events for Video REpresentation and Reasoning (CLEVRER)

The ability to reason about temporal and causal events from videos lies at the core of human intelligence. Most video reasoning benchmarks, however, focus on pattern recognition from complex visual and language input, instead of on causal structure. The study is focused on the complementary problem, exploring the temporal and causal structures behind videos of objects with simple visual appearance. CLEVRER is a diagnostic video dataset for systematic evaluation of computational models on a wide range of reasoning tasks. This dataset was introduced by the authors of the paper titled “CLEVRER: Collision Events for Video Representation and Reasoning.” [94] Motivated by the theory of human causal judgment, CLEVRER includes four types of question: descriptive (e.g., ‘what color’), explanatory (‘what’s responsible for’), predictive (‘what will happen next’), and counterfactual (‘what if’). Evaluation includes various state-of-the-art models for visual reasoning on a benchmark. While these models thrive on the perception-based task (descriptive), they perform poorly on the causal tasks (explanatory, predictive and counterfactual), suggesting that a principled approach for causal reasoning should incorporate the capability of both perceiving complex visual and language inputs, and understanding the underlying dynamics and causal relations.

Approach of various state-of-the-art visual reasoning models on CLEVRER

While these models perform well on descriptive questions, they lack the ability to perform causal reasoning and struggle on the explanatory, predictive, and counterfactual questions. We therefore identify three key elements that are essential to the task: recognition of the objects and events in the videos; modeling the dynamics and causal relations between the objects and events; and understanding of the symbolic logic behind the questions. As a first-step exploration of this principle, we study an oracle model, Neuro-Symbolic Dynamic Reasoning (NS-DR), that explicitly joins these components via a symbolic video representation and assess its performance and limitations.

Conclusion by evaluating team of CLEVRER with 4 types of questions

We present a systematic study of temporal and causal reasoning in videos. This profound and challenging problem deeply rooted to the fundamentals of human intelligence has just begun to be

studied with ‘modern’ AI tools. We introduce a set of benchmark tasks to better facilitate the research in this area. We also believe video understanding and reasoning should go beyond passive knowledge extraction and focus on building an internal understanding of the dynamics and causal relations, which is essential for practical applications such as dynamic robot manipulation under complex causal conditions. Our newly introduced CLEVRER dataset and the NS-DR model are preliminary steps toward this direction. We hope that with recent advances in graph networks, visual predictive models, and neuro-symbolic algorithms, the deep learning community can now revisit this classic problem in more realistic setups in the future, capturing true intelligence beyond pattern recognition.

Modular enhancements to neuro-symbolic dynamic reasoning model

Mao et al. [95] introduced the CLEVRER 1 dataset for systematic evaluation of computational models on descriptive, explanatory, predictive, and counterfactual questions about the movement of several objects with various shapes, colors, and materials. Noting that the state-of-the-art neural models had difficulty reasoning about temporal and causal structures for answering those questions, they proposed a neuro-symbolic model called NS-DR [95], which outperforms the previous models by using symbolic representation to allow for compositionality of vision, language, and dynamics. The result advocates that the use of explicit symbolic representation, combined with neural network perception, could significantly improve reasoning about complex visual events. On the other hand, this point is challenged by Ding et al. [96], who demonstrate that an end-to-end attention-based neural model with the right inductive bias could outperform NS-DR. Does this imply that neuro-symbolic models are inferior to end-to-end neural models for visual causal and temporal reasoning, contrary to what they were thought to be promising at? This note revisits the neuro-symbolic baseline model NS-DR. With the incorporation of more explicit causal and temporal constraints, we show that the enhanced model outperforms the previous models. This note briefly describes how we made modular improvements to NS-DR.

Our updates to NS-DR are relatively simple, thanks to its modular design. Without retraining the neural network models in NS-DR, the main reason for the improvement could be attributed to using explicit symbolic reasoning in ASP to determine what intermediate results the attention should be paid to and augmenting the mistakes in perception to follow physical

constraints. Our main improvement on counterfactual QA utilizes answer set programming (ASP) [97], a declarative logic programming paradigm that could encode various kinds of complex knowledge, including causal and temporal knowledge. For the CLEVRER task, we encode causal relationships among collision events.

CHAPTER 3 “COGENT WAY” KNOWLEDGE STRATEGY

“Cogent Way” (CW) project had been launched by Joint Staff Futures (JSF) J28 in early 2022. The main idea is to make data machine-interpretable, decision-ready, and actionable on par with how data is human-understandable and actionable. Without data architected for machine-interpretability/decision-readiness/actionability, none of the approaches described in this document as applied to AI/ML are possible. We believe this challenge requires utilization of Natural Language Processing (NLP) and using Deep Learning for training Natural Language (NL) Models combined with the Common- Sense Knowledge Graph (CSKG) introduced earlier in Chapter 2.

CW, at its core foundation, is focused on utilizing open knowledge standards based on Semantic Web Technologies approved and commercialized under World Wide Web Consortium (W3C) standard body. The crux of an interest to the CW project are semantic web ontologies and their integration with the Resource Definition Framework (RDF) and related standards (RDFS, RDF*, OWL, SPARQL/Geo-SPARQL, SWRL and others. CW project also sets as a major goal the integration of Labeled Property Graphs (LPG) with RDF/RDF* graphs being in-focus of the CW project. This includes a highly popular graph databases including Neo4j LPG graph database.

CHAPTER 4 NPS/NAVWAR KNOWLEDGE STRATEGY EXTENSION

Battlespace Situations and Context for Battlespace Situation

NPS team's focus was on "Naval Integration into Joint Data Strategies and Architectures in JADC2." Our team was predominantly focused on the Tactical Tier paying the greatest attention to the Tactical Edge. This critical part is, usually, excluded from the Enterprise System Architecture due to the organizations or standard bodies not having sufficient level of expertise at this critical organizational complexity where all warfighting engagement is, actually, taking place.

Therefore, Chapter 4 is devoted to what the real crux of the knowledge strategy at the tactical edge is. We believe that the amalgam of CW enabling infrastructure at the foundation and extension of CW for the tactical edge leads to a proper understanding on how to build the Enterprise for all tiers, starting from the higher-level tiers (e.g., Operational, Strategic/National) based on a Tactical Tier as an essential foundation. Tactical Tier's enabling infrastructure is a must for integrating this tier with higher-level Enterprise tiers. Tactical tier should not be excluded from the Enterprise Strategy developed by the CDOs from the DoD Services and Intel Agencies. Tactical Tier, due to Battlespace Situation-driven Context digitally born there, affects higher tiers.

Our NRP study of Naval Operational Architecture (NOA) back in 2019 has concluded that knowing knowledge at the tactical edge requires getting the knowledge of the situations. We get evidence that our initial determination back 2 years ago was visionary and insightful. This 2020 study informed the Navy on what truly matters and now, in 2022, be out to use for the benefit of JADC2. Our team believes following our recommendations JADC2 will be on its upward trajectory towards "joint" capabilities. In fact, we view "situation is a knowledge" main conclusion of 2019 study is a key insight for integration between the Navy and the JADC2 to reshape JADC2 data strategies and architectures while informed by the Navy. This would make it straightforward to integrate all warfighting organizations (i.e., DoD services, Coalition Partners) under JADC2 umbrella of data strategies and architectures.

During 2020 study we did not know of the existence of scientific approach capable of detecting novel situations. Rich Chase, a distinguished colleague from NAVWAR, as a prominent member of the team, made a startling discovery that nonmonotonic causal logic may track emerging threats by connecting threatening objects/platforms with "threatening effects" (i.e.,

missile types) and, through one or more causal links, via intervention DoLogic may compute counterfactual effects to defend against the threats. Such “causal detection of novel situations” may be applied to any adversarial actors with multiple level of causal effects in the Area of Interest (AOR). Such linkages are expressible by constructing and managing graphs depicting emerging situations. Each situation includes multiple “adversarial red objects” and “friendly C5ISR blue objects” with all objects interconnected via links that matter.

As a side note, it is worth mentioning that situations have to incorporate temporal-spatial exploitation cues from higher-level to tactical tier, including the tactical edge. In case of DCGS-N (Distributed Common Ground System - Navy), intel analysts recommend specific tactical AOR coordinates for consideration during the tactical edge exploitation. “Situation” concept offers persistent cuing of “exploitation tactical edge area of interest” between National and Tactical Tiers.

Now, when graph topology of a situation is known, what is also known is a low-level knowledge of all relations between “adversarial red” and “friendly blue objects”. Knowledge of a totality of relations in a situation is known as a context of a situation. Now, we possess a knowledge of not only high-level situations, but low-level context. Do situations exist during some interval of time? Yes, that is how we expect them to behave. Continuous monitoring may conclude of no causal relationship between emerging objects and multiple level effects-threats. This will identify an emergence of a new situation. In the meantime, existing situation may either continue in time or deprecate. It is important to stress that battlespace situations have geo-spatial shape ever changing in time. Such shapes may cross different AORs and must be carefully managed in defensive and offensive postures.

Management of situational context is a critical factor in scaling modern knowledge pipeline. After discussion of low-level context, it is essential to introduce high-level context. The example of high-level context is elevated to workflows managing situations requiring workflows to self-adapt to situations with dynamic context. This means that workflows gain prominence by enforcing on-the-fly sequencing by reaching high-reward world model states and situations thus accomplishing context-aware semi-optimal adaptation. This has been a long-term goal of DARPA from 2017 resulting in a 2-billion-dollar investment. As a consequence, the context management accomplishes reduction of contextual volume by summarizing contexts with expanding granularity making hierarchically and collaboratively integrated decision-makers getting cohesive.

Summarization of Situations and Situational Context

The question of “summarization” is triple-edged. What is an input for “summarization and what engine would be capable of doing “summarization”? How to ensure “input summarization streams” avoid duplication of knowledge by adhering to “knowledge in situ” principal? Finally, how to compute “summarization contextual output” to provide value to each role defined within JADC2? These questions will be addressed within the next sub-chapter we call “Summarization”.

We will start from a variety of inputs representing contexts for each causally monitored situation. First of all, it is a topology for each situation as it depicts unique context at a given timestamp. This reality means that we must embed knowledge of topologies of the situations as direct input to summarization engine to ensure there is no movement of topological data from point to point.

Secondly, we must summarize all internal and external inputs applied to the topology of the situation. “Internal inputs” are encoded into situation graph as functions which effect maneuvering movement at the edge. These functions are usually referred to as “actions”. They, usually, belong to one of two categories: moving and shooting. Relatively little attention is placed on “external inputs.” They are the events which more than anything else explains the volatility of the world model. The events are emanated from either operational environment (e.g., “blue forces are retreating due to realization they will be surrounded and made disabled), or from a variety of scientific environments which are either God-made (i.e., METOC) or man-made (i.e., EMW.)

Finally, “summarization engine” must be aware of all organizations in JADC2 and all organizational roles with each of the organizations. However, this represents organizations and people roles irrespectively of which organization and hierarchical sub-organizations and which hierarchically organized roles participate in any given mission. This means each situation and corresponding contexts must include a Mission ID. This makes total sense. When situations are detected, they must ensure there is a mapping between a situation and a mission situation is corresponding to. We arrived to a conclusion that support of JADC2 mission roles requires multi-dimensional cube with dimension hierarchies. This will provide a capability to associate “mission” dimension for any chosen combination of sub-organizations for specific group of roles in JADC2. Considering a need to summarize all artifacts (i.e., situations, workflows, objects, actions, events, etc.), a data structure of summarization engine must be capable to embed all representations within

a variety of artifacts making it a truly generic “information space” container. We concluded “summarization” engine requires to be a model with multi-dimensional and hierarchical capabilities within each of the dimensions. Container model be capable of imbedding 4D Cube Model for METOC, “situation” and “process graphs”. In addition, container model must embed entities, including objects, activities, events. Upon further examination, we may conclude container may require ontological embeddings to execute semantic logics-based reasoning.

Embracing 3rd Wave of AI: Contextual Adaptation

DARPA Information Innovation Office (I2O) in with visionary Dr. John Launchbury at the helm, in February 2017 proclaimed they are pursuing the path towards explanatory models by moving from Machine Learning/Deep Learning (ML/DL) to “Contextual Adaptation”. In fact, DARPA announced a transition 2nd “statistical” Wave of AI to 3rd “contextual” Wave of AI. We believe “Contextual Adaptation” is critically important to JADC2-based battlespace. In addition to understanding meaning of features in AI, “world model” should adapt to continuously evolving situations, characterized by situational contexts, in the joint battlespace.

The question to answer is how to express a need for a change from world model? First of all, any adaptation got to be contextual which implies to be mission centric. Secondly, world model is asking for relevant adaptation to its state via execution of a meaningful functional event. There are different manifestations of a functional event. It could be demanding to execute an action at the battlespace. Or react to a specific battlespace event from the operations environment. Finally, a change in one of the conditions in surrounding environment (i.e., METOC, EM, Acoustic, others) requires re-maneuvering to overcome emerging environmental threats in the physical world.

What is a perceived mechanism in the knowledge strategy to generate ad-hoc events based on the reasoning in accumulated knowledge for JADC2 operations? Cogent Way project is recommending establishing two ‘knowledge class layers’ within its knowledge strategy. “Contributing Knowledge Class” layer will act as landing zone in Extract-Load-Transform (ELT). This is where augmentation across different domains based on harmonization between diverse lexicons will be taking place. The 2nd “knowledge class layer” called “Aggregating knowledge system class” will be responsible for aggregation-by-linkage operation based on semantic web shallow-deep reasoning logic. Whenever deemed necessary, 2nd layer could be further extended by introducing the 3rd “deep reasoning knowledge class layer”. Such layer may be necessary for

managing battlespace in near real-time and recommending decisions which require other logics than causal inferencing logic. We are confident “automated augmentation” with, potentially, human-on-the-loop requires reasoning stack as opposed to viewing this task as a multi-sensor fusion. The goal of reasoning layers is to generate operational ad-hoc events to send back to the “Contributing Knowledge Class” via a feedback loop, and then, to Producer authoritative systems.

Upon concluding of a need of in-time generated events from the “Aggregating Knowledge Systems”, there is a need to decide where such events should be distributed to. Our notion is that discussed events must be streamed, through a feedback loop, to the “Contributing Knowledge Systems” to let the latter type of the Knowledge Graph (KG) group to be responsible to inform authoritative systems of their need to adapt mission contexts based on the type and reification of real-time ad-hoc events generated by the “Aggregation Knowledge Systems.” Described single or collaborative workflow executed from a feedback loop for any combination of domains/sub-domains is a must to have to accomplish real-time contextual adaptation.

Moving from Metadata Catalog to Contextual Data Catalog

Two members of our team attended JADC2 Architecture Symposium back in March 2022 in Suffolk, Virginia. The biggest impression from this high-level event was hearing about high stakes JADC2 has placed on a “Data Fabric”. The concept of a data fabric has been defined as “a DoD federated data environment for sharing information through interfaces and services to discover, understand and exchange data with partners across all domains, security levels and echelons.” [98]

Being knowledge-understanding-wisdom-and-beyond centric, our team feels JADC2 is not the right project where DoD should be experimenting with Data-Information quagmire. We, simply, have no time to be in the “coal mines”. Our goal is to build “skyscrapers”. We should be well on the ground and as high as we could be. We should not perfect “coal mines”. They are not going to help us to produce another kind of energy. There is lot of comparison between data and oil. We hear expression “data is a new oil”. We should not focus on data, we should climb a ladder to “knowledge”, “understanding”, “wisdom” and beyond. Should we rather move from coal mining onto gold mining? If, according to RADM Danelle Barrett, “Data is Oil”, then our team proclaims that “Knowledge is Gold.”



Figure 2 Data is Coal and Knowledge is Gold

This moves us to a discussion if Metadata Catalog is so robust to be carried into the future. According to a Data Fabric presentation, everything there rotates around the Metadata Catalog. It is like our galactic with planets rotating around the Sun. Is it really so “bullet proof” to be a foundation for all tiers of the Enterprise including the Tactical Edge?

Stepping into the World of Contextual Data Catalog

One example could easily destroy a myth of a power of Metadata Catalog. Is it contextually aware? If we have tracks passing by in a close vicinity from each other in 4-D area, are they focused on the same mission, or just decided to take a short route over deconflicted areas to get faster to their point of destination. They could be humans, or ordnance, including various missile types.

Are these objects intervening in conducting the mission? Of course, they do. Fire Support Coordinators (FSC) just finished the deconfliction. Surprise, surprise, Somebody is flying over. Are they red, or blue, or white? We need to follow the rules of engagement to avoid fratricide. All the calculation for deconflicting the space must start all over again. Blue assets were not supposed to be here! This is what knowledge of context helps with. And Metadata Catalogs cannot do that!



Figure 3 Metadata Catalog (single track) vs Contextual Data Catalog

What is the difference between the images on the left and right? The left image depicts a flight of space station in the outer space. There is no background or, in general, context to relate to. All we see is a track made by a space station. We could identify such station as it has visible features. The image on the right is different. In addition to two warriors, male and female, there are visible surrounding objects and landscapes. One may see large white stone and even larger white tock. There is a tall grass and flowers with one of them, on the left, being red. Finally, behind the large field of grass (or, potentially, some grain plants), there are dark green woods. All of these additions to two warriors depict the context. In general, context could be meaningful relations between objects. Battlespace object relations is a good representation of a context, but it's not all.

CHAPTER 5 CONCLUSIONS

This study has focused on significance of a knowledge strategy for not only the Naval Operational Architecture, including continuous efforts by Overmatch, but for the success of the whole of the JADC2. The main conclusion of our study is that Joint Missions must have a Knowledge Strategy at its foundation. The best way Knowledge Strategy shows its significance is for JADC2. In a sense, Knowledge Strategy and JADC2 have a symbiotic relationship.

The reason it's the case is because acting as a Joint Force requires a knowledge of validated data. Attempt to lead forces in a Joint mode without a knowledge of a data is like blind walking in the park. It's critical to have a shared understanding of what knowledge is. Specifically, knowledge of what. Our study, by applying abstractions of knowledge and associated context, declared a thesis 'knowledge' of situations is what is paramount in kinetic and non-kinetic battlespaces. This makes 'context' a knowledge of significant relations of subjects-objects participating in a given situation.

The study made it quite certain that causality based on causal diagram developed by the SMEs is a key to understanding situation and corresponding context. Due to situations and contexts changing continuously in a rapid fashion, causal inferencing must also run continuously in the background. Our team is also making an argument that causal and other logics related to situations/contexts, movements and events must be all non-monotonic. A logic is non-monotonic if some conclusions can be invalidated by adding more knowledge.

Our further conclusion is that for data to act as an input to knowledge, all data must have a property of "mergeable". Otherwise, a data piece that cannot be merged with other pieces of data can not affect eventual knowledge derived from the merged data. The same applies to "knowledge" as it also must have a "mergeable" property. Otherwise, it will not be possible to create knowledge out of knowledge which is a must to do in a role hierarchy. There is a need for two powerful operators executed in one step: "Merge and Summarize". The reason why Knowledge Graphs (KGs) are adopted by knowledge projects like Cogent Way is due to a power that triplets in a given Knowledge Graph (KG) are mergeable with some of the triples in other KGs. This will be always true in triplets are guaranteed not to be lossy. KG algorithms ensure none of the algorithms won't exhibit "loss of data" algorithmic property.

LIST OF ACRONYMS

AI	Artificial Intelligence
AI/ML	Artificial Intelligence / Machine Learning
AOR(s)	Area(s) Of Responsibility
ASP (1 st acronym)	Answer Set Prolog
ASP (2 nd acronym)	Answer Set Programming Paradigm
C5ISR	Command, Control, Computers, Communications, Cyber, Intelligence, Surveillance, Reconnaissance
CCALC	Causal CALCulator
CSKG	CommonSense Knowledge Graph
CW	Cogent Way project initiated by Joint Staff Futures
DAG	Direct Acyclic Graph
DARPA	Defense Advanced Research Projects Agency
DCGS-N	Distributed Command Ground System – Navy
ELT	Extract, Load, Transform
EMW	ElectroMagnetic Warfare
GeoSPARQL	Interface to support representing and querying geospatial data on the semantic web
JADC2	Joint All-Domain Command and Control
JS J6	Joint Staff Joint Command & Control Branch
KG(s)	Knowledge Graph(s)
KRR	Knowledge Representation and Reasoning
LMs	Language Models
LPG	Labeled Property Graph
METOC	Meteorology and Oceanography
ML	Machine Learning

ML/DL	Machine Learning / Deep Learning
MLN	Markov Logic Network
NAVWAR	Naval Information Warfare Systems Command
Neo4j	Graph DBMS developed by Neo4j
NGA	National Geospatial-Intelligence Agency
NRO	National Reconnaissance Office
NSA	National Security Agency
NS-DR	Neuro-Symbolic Dynamic Reasoning
OOP	Object-Oriented Programming
OWL	Web Ontology Language
Prolog	Programming Language frequently used within AI
QA	Query Answering
QSR	Qualitative Spatial Reasoning
RDF	Resource Description Framework
RDFS	Resource Description Framework Schema
RDF-star or RDF*	RDF extension
SA	Situation Awareness
SOF	Special Operations Forces
SPARQL	SPARQL Protocol And RDF Query Language
SWRL	Semantic Web Rule Language

LIST OF REFERENCES

- [1] Katz, J.: “Navy advances ‘operational architecture’ for storing, analyzing, moving data”, www.insidedefense.com, September 3, 2020
- [2] John F. Sowa, Knowledge Representation: Logical, Philosophical, and Computational Foundations, Brooks/Cole, 2000
- [3] Stuart Russell, Peter Norvig, 4th edition, *Artificial Intelligence: A Modern Approach*. Published by Pearson (April 28, 2020)
- [4] Frame Logic (aka F-logic), <https://en.m.wikipedia.org/wiki/F-logic>
- [5] Gerhard, Brevka; Ilkka, Niemelä; Mirosław, Truszczyński; Chapter 6, “Nonmonotonic Reasoning”, pp. 268; Edited by Van Harmelen, Frank; Lifschitz, Vladimir; Porter, Bruce; Lierler, Yuliya, “Handbook on Knowledge Representation” (2008), Faculty Books and Monographs
- [6] Gerhard, Brevka; Ilkka, Niemelä; Mirosław, Truszczyński; Chapter 6, “Nonmonotonic Reasoning”, pages 239-240; Van Harmelen, Frank; Lifschitz, Vladimir; Porter, Bruce; Lierler, Yuliya, “Handbook on Knowledge Representation” (2008), Faculty Books and Monographs
- [7] M. Gelfond and V. Lifschitz. The stable model semantics for logic programming. In *Proceedings of ICLP-88*, pages 1070–1080, 1988.
- [8] M. Gelfond and V. Lifschitz. Classical negation in logic programs and disjunctive databases. *New Generation Computing*, 9(3–4):365–386, 1991.
- [9] R. Kowalski. *Logic for Problem Solving*. North-Holland, 1979.
- [10] J. McCarthy. In V. Lifschitz, editor. *Formalization of Common Sense*. Ablex, 1990.
- [11] K. Clark. Negation as failure. In *Logic and Data Bases*, pages 293–322. Plenum Press, 1978.
- [12] J. Minker. On indefinite data bases and the closed world assumption. In *Proceedings of CADE-82*, pages 292–308, 1982.
- [13] J. Minker. Logic and databases: a 20-year retrospective. In H. Levesque and F. Pirri, editors. *Logical Foundations for Cognitive Agents: Contributions in Honor of Ray Reiter*, pages 234–299. Springer, 1999.
- [14] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13(1–2):81–132, 1980.
- [15] R.C. Moore. Semantical considerations on nonmonotonic logic. In *Proceedings of the 8th International Joint Conference on Artificial Intelligence*, pages 272–279. Morgan Kaufmann, August 1983.
- [16] V.W. Marek and M. Truszczyński. *Nonmonotonic Logics; Context Dependent Reasoning*. Springer-Verlag, Berlin, 1993.
- [17] I. Niemela. Logic programs with stable model semantics as a constraint programming paradigm. *Annals of Mathematics and Artificial Intelligence*, 25(3–4):241–247, 1999.
- [18] V.W. Marek and M. Truszczyński. Stable models and an alternative logic programming paradigm. In *The Logic Programming Paradigm: a 25-Year Perspective*, pages 375–398. Springer-Verlag, Berlin, 1999.
- [19] M. Nogueira, M. Balduccini, M. Gelfond, R. Watson, and M. Barry. An A-Prolog decision support system for the Space Shuttle. In *PADL 2001*, pages 169–183, 2001.
- [20] T. Soininen and I. Niemela. Developing a declarative rule language for applications in product configuration. In *Proceedings of International Symposium on Practical Aspects of Declarative Languages*, pages 305–319, 1998.

- [21] C. Baral, K. Chancellor, N. Tran, A. Joy, and M. Berens. A knowledge based approach for representing and reasoning about cell signaling networks. In *Proceedings of European Conference on Computational Biology, Supplement on Bioinformatics*, pages 15–22, 2004.
- [22] D.R. Brooks, E. Erdem, J.W. Minett, and D. Ringe. Character-based cladistics and answer set programming. In *Proceedings of International Symposium on Practical Aspects of Declarative Languages*, pages 37–51, 2005.
- [23] W. Faber, N. Leone, and G. Pfeifer. Recursive aggregates in disjunctive logic programs: Semantics and complexity. In *Proceedings of the 8th European Conference on Artificial Intelligence (JELIA 2004)*, pages 200–212, 2004.
- [24] W. Faber. Unfounded sets for disjunctive logic programs with arbitrary aggregates. In *Proc. of 8th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR 2005)*, *LNAI*, vol. 3662, pages 40–52. Springer-Verlag, Berlin, 2005.
- [25] T.C. Son and E. Pontelli. A constructive semantic characterization of aggregates in answer set programming. *Theory and Practice of Logic Programming*, 7:355–375, 2007.
- [26] N. Pelov, M. Denecker, and M. Bruynooghe. Well-founded and stable semantics of logic programs with aggregates. *Theory and Practice of Logic Programming*, 7:355–375, 2007.
- [27] P. Ferraris. Answer sets for propositional theories. In *Proceedings of International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR)*, pages 119–131, 2005.
- [28] P. Ferraris and V. Lifschitz. Weight constraints as nested expressions. *Theory and Practice of Logic Programming*, 5:45–74, 2005.
- [29] G. Brewka. Logic programming with ordered disjunction. In *Proceedings of AAAI-02*, pages 100–105, 2002.
- [30] J.P. Delgrande, T. Schaub, H. Tompits, and K. Wang. A classification and survey of preference handling approaches in nonmonotonic reasoning. *Computational Intelligence*, 20(2):308–334, 2004.
- [31] C. Sakama and K. Inoue. Prioritized logic programming and its application to commonsense reasoning. *Artificial Intelligence*, 123:185–222, 2000.
- [32] F. Buccafurri, N. Leone, and P. Rullo. Adding weak constraints to disjunctive datalog. In *Proceedings of the 1997 Joint Conference on Declarative Programming APPIA-GULP-PRODE’97*, 1997.
- [33] M. Balduccini and M. Gelfond. Logic programs with consistency-restoring rules. In P. Doherty, J. McCarthy, and M.-A. Williams, editors, *International Symposium on Logical Formalization of Commonsense Reasoning, AAAI 2003 Spring Symposium Series*, pages 9–18, March 2003.
- [34] C. Baral, M. Gelfond, and N. Rushton. Probabilistic reasoning with answer sets. In *Proceedings of LPNMR-7*, January 2004.
- [35] G. Brewka. Answer sets: From constraint programming towards qualitative optimization. In *Proc. of 7th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR-04)*, pages 34–46. Springer-Verlag, Berlin, 2004.
- [36] P.J. Hayes and J. McCarthy. Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer and D. Michie, editors. *Machine Intelligence 4*, pages 463–502. Edinburgh University Press, 1969.

- [37] M. Shanahan. *Solving the Frame Problem: A Mathematical Investigation of the Commonsense Law of Inertia*. MIT Press, 1997.
- [38] P.J. Hayes and J. McCarthy. Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer and D. Michie, editors. *Machine Intelligence 4*, pages 463–502. Edinburgh University Press, 1969.
- [39] R. Reiter. *Knowledge in Action—Logical Foundations for Specifying and Implementing Dynamical Systems*. MIT Press, September 2001.
- [40] R.A. Kowalski and M. Sergot. A logic-based calculus of events. *New Generation Computing*, 4(4):319–340, 1986.
- [41] E.T. Mueller. *Commonsense Reasoning*. Morgan Kaufmann, 2006.
- [42] C. Baral. *Knowledge Representation, Reasoning and Declarative Problem Solving with Answer Sets*. Cambridge University Press, 2003.
- [43] M. Gelfond and N. Leone. Knowledge representation and logic programming. *Artificial Intelligence*, 138(1–2), 2002.
- [44] Michael Gelfond; Chapter 7, “Answer Sets”, pages 285-286, page 310; Edited by Van Harmelen, Frank; Lifschitz, Vladimir; Porter, Bruce; Lierler, Yuliya, “Handbook on Knowledge Representation” (2008), *Faculty Books and Monographs*
- [45] J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, New York, 2000.
- [46] Chapter 11, “Bayesian Networks”, page 467, page 497-498; Edited by Van Harmelen, Frank; Lifschitz, Vladimir; Porter, Bruce; Lierler, Yuliya, “Handbook on Knowledge Representation” (2008), *Faculty Books and Monographs*
- [47] P. Doherty, J. Kvarnström, Chapter 18, “Temporal Action Logic”, 2007, page 709; Edited by Van Harmelen, Frank; Lifschitz, Vladimir; Porter, Bruce; Lierler, Yuliya, “Handbook on Knowledge Representation” (2008), *Faculty Books and Monographs*
- [48] F. Lin, Chapter 16, “Situation Calculus”, 2007, p. 649; Edited by Van Harmelen, Frank; Lifschitz, Vladimir; Porter, Bruce; Lierler, Yuliya, “Handbook on Knowledge Representation” (2008), *Faculty Books and Monographs*
- [49] E. Mueller, Chapter 17, “Event Calculus”, pages 700-701; Edited by Van Harmelen, Frank; Lifschitz, Vladimir; Porter, Bruce; Lierler, Yuliya, “Handbook on Knowledge Representation” (2008), *Faculty Books and Monographs*
- [50] A. G. Cohn, J. Renz, Chapter 13, “Qualitative Spatial Representation and Reasoning”, p. 552; Edited by Van Harmelen, Frank; Lifschitz, Vladimir; Porter, Bruce; Lierler, Yuliya, “Handbook on Knowledge Representation” (2008), *Faculty Books and Monographs*
- [51] F. Rossi, P. Van Beek, T. Walsh, Chapter 4, “Constraint Programming”, p. 181; Edited by Van Harmelen, Frank; Lifschitz, Vladimir; Porter, Bruce; Lierler, Yuliya, “Handbook on Knowledge Representation” (2008), *Faculty Books and Monographs*
- [52] P. Traverso. Chapter 22, “Automated Planning”, p. 841; Edited by Van Harmelen, Frank; Lifschitz, Vladimir; Porter, Bruce; Lierler, Yuliya, “Handbook on Knowledge Representation” (2008), *Faculty Books and Monographs*
- [53] K. Forbus. Chapter 9, “Qualitative Modeling”, p. 361; Edited by Van Harmelen, Frank; Lifschitz, Vladimir; Porter, Bruce; Lierler, Yuliya, “Handbook on Knowledge Representation” (2008), *Faculty Books and Monographs*

- [54] F. Baader, I. Horrocks, U. Sattler, Chapter 3, “Description Logics”, p. 137, *Edited by Van Harmelen, Frank; Lifschitz, Vladimir; Porter, Bruce; Lierler, Yuliya, “Handbook on Knowledge Representation” (2008), Faculty Books and Monographs*
- [55] Y. Moses, Chapter 15, “Reasoning about Knowledge and Belief”, p. 621, *Edited by Van Harmelen, Frank; Lifschitz, Vladimir; Porter, Bruce; Lierler, Yuliya, “Handbook on Knowledge Representation” (2008), Faculty Books and Monographs*
- [56] C. Baral, M. Gelfond, and N. Rushton. Probabilistic reasoning with answer sets. In *Proceedings of LPNMR-7*, pages 21–33, Jan 2004.
- [57] W. van der Hoek and M. Wooldridge, Chapter 24, “Multi-Agent System”, pp. 887-888, *Edited by Van Harmelen, Frank; Lifschitz, Vladimir; Porter, Bruce; Lierler, Yuliya, “Handbook on Knowledge Representation” (2008), Faculty Books and Monographs*
- [58] M. Fisher, D. Gabbay, and L. Vila, editors. *Handbook of Temporal Reasoning in Artificial Intelligence, Foundations of Artificial Intelligence*, vol. 1. Elsevier Press, 2005.
- [59] D. Gabbay, I. Hodkinson, and M. Reynolds. *Temporal Logic: Mathematical Foundations and Computational Aspects*, vol. 1. Clarendon Press, Oxford, 1994.
- [60] L. Bolc and A. Szalas, editors. *Time and Logic: A Computational Approach*. Univ. College London Press, 1995.
- [61] P. Gribomont and P. Wolper. In *From Modal Logic to Deductive Databases: Introducing a Logic Based Approach to Artificial Intelligence*, pages 165–234. Wiley, 1989 (chapter *Temporal Logic*).
- [62] D. Gabbay, M. Reynolds, and M. Finger. *Temporal Logic: Mathematical Foundations and Computational Aspects*, vol. 2. Clarendon Press, Oxford, 2000.
- [63] I. Hodkinson and M. Reynolds. Temporal logic. In P. Blackburn, J. van Benthem, and F. Wolter, editors. *Handbook of Modal Logic*. Elsevier, 2006 (Chapter 11).
- [64] M. Fischer. Chapter 12, “Temporal Representation and Reasoning”, p. 513; *Edited by Van Harmelen, Frank; Lifschitz, Vladimir; Porter, Bruce; Lierler, Yuliya, “Handbook on Knowledge Representation” (2008), Faculty Books and Monographs*
- [65] K.D. Forbus, P. Nielsen, and B. Faltings. Qualitative spatial reasoning: The clock project. *Artificial Intelligence*, 51:417–471, 1991.
- [66] D.S. Weld and J. De Kleer, editors. *Readings in Qualitative Reasoning About Physical Systems*. Morgan Kaufman, San Mateo, CA, 1990.
- [67] R. Fagin, J.Y. Halpern, Y. Moses, and M.Y. Vardi. *Reasoning about Knowledge*. The MIT Press, 2003.
- [68] J. McCarthy. Situations, actions, and causal laws. In M. Minsky, editor. *Semantic Information Processing*, pages 410–417. MIT Press, Cambridge, MA, 1968.
- [69] J. McCarthy and P. Hayes. Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer and D. Michie, editors. *Machine Intelligence*, vol. 4, pages 463–502. Edinburgh University Press, Edinburgh, 1969.
- [70] R. Reiter. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. The MIT Press, 2001.
- [71] D. Harel. *First-Order Dynamic Logic. Lecture Notes in Computer Science*, vol. 68. Springer-Verlag, New York, 1979.
- [72] R.A. Kowalski and M.J. Sergot. A logic-based calculus of events. *New Generation Computing*, 4(1):67–95, 1986.

- [73] R. Miller and M. Shanahan. Some alternative formulations of the event calculus. In A.C. Kakas and F. Sadri, editors. *Computational Logic: Logic Programming and Beyond: Essays in Honour of Robert A. Kowalski, Part II, Lecture Notes in Computer Science*, vol. 2408, pages 452–490. Springer, Berlin, 2002.
- [74] E.T. Mueller. *Commonsense Reasoning*. Morgan Kaufmann, San Francisco, 2006.
- [75] M. Shanahan. The event calculus explained. In M.J. Wooldridge and M.M. Veloso, editors. *Artificial Intelligence Today: Recent Trends and Developments, Lecture Notes in Computer Science*, vol. 1600, pages 409–430. Springer, Berlin, 1999.
- [76] R. Miller and M. Shanahan. Narratives in the situation calculus. *Journal of Logic and Computation*, 4(5):513–530, 1994.
- [77] Y. Shoham. *Reasoning about Action and Change*. MIT Press, 1987.
- [78] R. Reiter. *Knowledge in Action*. MIT Press, 2001.
- [79] E. Sandewall. *Features and Fluents: A Systematic Approach to the Representation of Knowledge about Dynamical Systems*, vol. 1. Oxford University Press, 1994.
- [80] M. Thielscher. *Reasoning Robots—The Art and Science of Programming Robotic Agents*. Springer, 2005.
- [81] C. Baral and M. Gelfond. Logic programming and reasoning about actions. In M. Fisher, D.M. Gabbay, and L. Vila, editors. *Handbook of Temporal Reasoning in Artificial Intelligence*. Elsevier Publications, 2005.
- [82] J. McCarthy. Programs with common sense. In *Proceedings of the Teddington Conference on the Mechanization of Thought Processes*, London, pages 75–91. Her Majesty’s Stationary Office, 1959.
- [83] J. McCarthy. *Formalization of Common Sense, papers by John McCarthy edited by Vladimir Lifschitz*. Ablex, 1990.
- [84] Y. Shoham. Nonmonotonic logics: Meaning and utility. In J.P. McDermott, editor. *Proceedings of the Tenth International Joint Conference on Artificial Intelligence (IJCAI-1987)*, Milan, Italy, August 1987, pages 388–393. Morgan Kaufmann Publishers, Los Altos, CA, USA, 1987.
- [85] N. McCain and H. Turner. Causal theories of action and change. In *Proc. of AAAI-97*, pages 460–465, 1997.
- [86] F. Lin. Embracing causality in specifying the indirect effects of actions. In *Proc. of IJCAI-95*, pages 1985–1991, 1995.
- [87] N. McCain and H. Turner. A causal theory of ramifications and qualifications. In *Proc. of IJCAI-95*, pages 1978–1984, 1995.
- [88] H. Turner. Representing actions in logic programs and default theories: A situation calculus approach. *Journal of Logic Programming*, 31(1–3):245–298, 1997.
- [89] E. Giunchiglia, J. Lee, V. Lifschitz, N. McCain, and H. Turner. Nonmonotonic causal theories. *Artificial Intelligence*, 153(1–2):49–104, 2004.
- [90] N. McCain. Causality in commonsense reasoning about actions. PhD Dissertation, University of Texas at Austin, Department of Computer Sciences, 1997.
- [91] Research interests of Dr. Farough Arabshahi, Metadata (former company name: Facebook), <https://forougha.github.io/research.html>
- [92] Forough Arabshahi, “Neuro-Symbolic Algorithms for Automated Reasoning”, Allen Institute for AI, May 27, 2020, <https://www.youtube.com/watch?v=83sTGeR6kdg>

- [93] COMET-ATOMIC 2020: On Symbolic and Neural Commonsense Knowledge Graphs, Jena D. Hwang, Chandra Bhagavatula, Ronan Le Bras, Jeff Da, Keisuke Sakaguchi, Antoine Bosselut, Yejin Choi. Published in AAAI 12 October 2020 Computer Science
- [94] CLEVRER: “Collision Events for Video Representation and Reasoning”, Kexin Yi*, Chuang Gan*, Yunzhu Li, Pushmeet Kohli, Jiajun Wu, Antonio Torralba, Joshua B. Tenenbaum, ICLR, 2020
- [95] Jiayuan Mao, Chuang Gan, Pushmeet Kohli, Joshua B. Tenenbaum, and Jiajun Wu. The neuro-symbolic concept learner: interpreting scenes, words, and sentences from natural supervision. In ICLR, 2019.
- [96] David Ding, Felix Hill, Adam Santoro, Malcolm Reynolds, and Matt Botvinick. Attention over learned object embeddings enables complex visual reasoning. arXiv 2012.08508, 2020.
- [97] Vladimir Lifschitz. What is answer set programming? In Proceedings of the AAAI Conference on Artificial Intelligence, pages 1594–1597. MIT Press, 2008.
- [98] “Data Takes a Central Role in JADC2 Strategy.” By Kimberly Underwood, Feb 23, 2021, <https://www.afcea.org/signal-media/data-takes-central-role-jadc2-strategy> , Signal-Media
- [99] Pedro Domingos, “The Master Algorithm: How the Quest for the Ultimate Learning Machine Will Remake Our World”, Basic Books, September 22, 2015

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
Ft. Belvoir, Virginia
2. Dudley Knox Library
Naval Postgraduate School
Monterey, California
3. Research Sponsored Programs Office, Code 41
Naval Postgraduate School
Monterey, California

THIS PAGE INTENTIONALLY LEFT BLANK